they could point to each one.

One day, of course, BBN increased the size of the memory, they doubled it to 32-K. And they forgot to add more pointers. When we did our measurements, all the traffic came through our IMP, creating a large number of small packets, which required a large number of pointers and they didn't have enough. So, data was being put into storage and (Inaudible) point to to ever release the storage. (Claps Hands) The network came down. Christmas deadlock. That was one of them. There was a large number of others too.

QUESTION

LEN KLEINROCK

The relationship between Larry Roberts and me is a very special one. We engaged in a large number of, uh, interesting ventures and

I'm a graduate student, running my difficult simulation at Lincoln Laboratory, of these computer networks. I had the midnight to seven a.m. shift four days a week and they were not contiguous days, so you can i

adventures, both as students and as professionals afterwards. I can tell you, for example, at least three stories. Story number one.

I'm a graduate student, running my difficult simulation at Lincoln Laboratory, of these computer networks. I had the midnight to seven a.m. shift four days a week and they were not contiguous days, so you can imagine what that did to my sleeping schedule. One night, late at night, about two in the morning, I was there in front of this machine, the TX-2 on which I was running the simulations. The TX-2 is the, one of the first experimental transistorized computers, a magnificent state of the art machine. Experimental machine, so typically, pieces of the machine were in repair or under modification at any one time.

And typically the console had empty holes in it. There I am, two in the morning, late

at night, tired as a dog, all alone and this machine. And you know, when you're alone with the machine at night, you listen to all the sounds. And if any sound is wrong, you pick it up immediately, (Snaps Fingers), and you're worried you're gonna break something. I'm working there and all of a sudden I hear, (Makes Hissing Sound), a strange windy sound, and my heart jumped, cause I knew something was about to go wrong, and I couldn't find out what it was, so I, (Makes Hissing Sound), I was looking around, I'm, finally my eyes scan the console and there in an empty space, I saw two eyes looking at me.

It was the son of a bitch, Larry Roberts. He was back there, just frightening me, and he successfully did it. Now did I say, is son of a bitch?

(OFF MIKE)

QUESTION

LEN KLEINROCK

There's lots more, okay. Another story with Larry Roberts is, one day he came to me, this was again 1968, in the heyday of the ARPANET not yet built, and he pointed out to me that silver coins, a quarter, made out of silver, was worth more melted than it was as a twenty-five cent piece. But it was illegal to melt the coins. So Larry and I got together and said, listen, we're gonna collect these coins, for the day when we can melt them.

So I invented a machine that could sort quarters. Cause at that time, in '68, there were both mixed silver and these clad quarters, which were worth only twenty-five cents as coin, not melted. The trick is to separate them. I built a machine that could handle ten thousand dollars an hour. You feed it in, a magnetic field separates it two, and I had a bayonet blade which define(?) the final separation in my garage.

Quarters are flying all over the place.

I then had to go and find quarters. I went to the banks, they wouldn't sell them to me, because the U.S. government was collecting them. I went to the vending machines, ended up talking to the Mafia, I couldn't get them there. I finally had to describe my machine and buy them from some sleazy looking people who were using Mexican labor to hand pick the quarters.

So Larry and I got together with two other people that, one was Barry Wessler(?), Larry's colleague at ARPA, and one was Walter Karplus(?), a colleague of mine at UCLA, and we put some money together. And here was the deal. Walter was the money man, he had five thousand dollars, we had nothing. You take the five thousand dollars, you buy five thousand dollars worth of silver quarters. You bring it to the bank, tell the bank, I want to borrow five thousand dollars, and for collateral I'm gonna

give you five thousand, put it in storage for me.

But they're gonna charge you interest. You get the five thousand, you buy another five thousand of quarters. You keep this until the premium you paid for the coin and the interest you're paying the bank has you up to here. So the deal was this. Larry decided to go for ten thousand, Barry for ten thousand, Walter for ten thousand, and me for a hundred thousand. I said, I'm gonna do it, I'm gonna do it right. And it all happened.

Collected it all. The price started dropping. When it was low enough, the government said it's legal to melt. The price started rising slowly and all my colleagues were very unhappy that they were losing money. When it reached a point where they could break even, most of them sold out. I kept in. Finally when it went up higher, I sold out and I made a fortune in silver, I still have bags of silver in a

safe deposit box right now, worth a lot of money.

QUESTION

LEN KLEINROCK

Okay. Blackjack.

QUESTION

LEN KLEINROCK

Okay. Larry and I were also avid blackjack players. Because between the two of us, we had devised a system around the same time that Philip(?) had devised his system, where you could beat blackjack. You know, blackjack is (Inaudible) with memory, if you know what's going on, you count the cards, you can beat the heck out of it.

Well, Larry and I ventured to Las Vegas more than once, and we started cleaning up. So much so that we got thrown out of clubs a few times. Larry also recognized that we could beat roulette. If the roulette wheel, you know,

roulette is a very, very precise game. The ball rotates at a certain speed, Newton's laws of motion say it's gonna fall off at a certain point. If you can predict which half it's gonna fall off at, you ... you can make a fortune. The trick is, to measure the speed of the ball which is going one way, the speed of the wheel going the other way, measure the, slow down the velocity and predict it all.

So the net result is, one could build a little machine that would make that work. But we needed to get some data as to how the wheel and how fast the ball were going. So Larry took a tape recorder, put the microphone in his hand, wrapped it up like he had a broken arm, put the tape recorder in his jacket and he and ... and I went down to the roulette table. He put his hand next to the roulette wheel, so it could hear the ball going (Makes Whirring Engine Noises), and my job as to play roulette to distract

them.

So we went down there, we were working fine except for one flaw. I started winning. (Laughs) And it drew attention to me. And the pit boss sees I'm winning, and sees this guy with a broken arm with his hand near the roulette wheel, so he grabs Larry's broken arm and he says, let me see your arm. And it's a broken arm, right? Larry and I made, uh, fast tracks out of there that day.

(OFF MIKE)

QUESTION

(CUT)

(OFF MIKE)

QUESTION

LEN KLEINROCK

The ARPANET experience we're describing right now was something that happened beginning in 1967, '69 and those early days, which were really very exciting days, I mean, you have to

understand, we were creating a new technology. We didn't realize the full impact of it, but the idea of creating this networking capability and seeing the applications and seeing the growth, and working with these people, all of whom were stars, was an absolutely unique experience.

Uh, you know, one of the high points of one's career. The same time, I was busy generating Ph.D. students who were experts in this area. One of my greatest outputs are the thirty-five Ph.D. students I've graduated, who populate the world and carry that expertise forward. But moving us into today's world, you know, we're in the middle of a situation where ... where three critical events have juxtaposed to put us at a very important juncture in the field of networking and information technology.

(TAPE SPEED VARIES)

LEN KLEINROCK

First even it, the Internet has exploded. Second

event is, this administration is promoting the information superhighway in a very powerful way. The third event is, that the entertainment, the telephone and the cable TV industry have recognized there is a big commercial market out there. And they're ready to pour fifty billion dollars a pop into providing this information superhighway.

Now the problem is, where's the guidance for all of this? Where is the technology framework to help this thing move forward? There are various visions, but where is the meat on those visions? Two months ago, in May, I just, I chaired a committee for the National Research Council, which monitors basically National Academy of Science and Engineering, and we produced a report which came out in May, called "Realizing the Information Future", and the sub-title is, "The ... The Internet and Beyond". Or is it "Beyond the Internet"? I

should do that again, shouldn't I?

(OFF MIKE)

QUESTION

LEN KLEINROCK

This report just came out ...

QUESTION

LEN KLEINROCK

This report just came out, and it's basically become a landmark report. This committee of sixteen people, chaired by myself, I'm going around with a couple of committee members, basically briefing the chief scientists, Jack Gibbons(?), the head of NIST(?), Office of Management and Budget, Department of Commerce, we did a Congressional sub-committee briefing as well, promoting this, these ideas that we have.

And what we've done is we've created a technology framework for this vision, which explains the key architectural issues one has to

consider in bringing this superhighway about. It also addresses things like the needs of the public interest groups like education, libraries, science. Addresses the financial issues, how you fund and charge for this thing. And more importantly, what the role of government should be.

So, we're right now at a very critical time when a little bit of guidance can go a long way in making sure that we don't end up simply with 500 channels of TV and an RJ-11(?) jack for a back channel. We need a true symmetric(?) (Inaudible) service which provides the flexibility and the growth capability. That's an activity right now that a number of people in the community, Bob Kahn is on this committee, by the way, as well. But we, it's very important that we get this right this time, with all these billions about to be invested, and it's the private sector that's gonna invest the money, not the government.

So the trick is how to, how to encourage the private sector to have a common view that serves them all. Let's look forward to the future now. Where is this gonna take us in another twenty-five years? Well, first of all, let's remember this. Twenty-five years ago, there were some very significant events that happened. One is we put a man on the moon, on July twentieth. Twenty-four billion dollars and eight years of work. The other one was Woodstock. Another one was Charles Manson. Today, August tenth, twenty-five years ago. And the fourth, of course, was the ARPANET.

The ARPANET was an inexpensive experiment, high risk, high pay off, that the government invested in, and the impact is clearly enormous. Twenty-five years from now, we'll be seeing (Inaudible) like that, where we'll look back and say, we had a time here when something critical was happening and I think the answer is

yes. I think the notion of this overall vision for this network which now leaves the science community, enters the commercial arena, the consumer, education, public interest, is gonna be a real juncture in the way things take place.

At the same time, there's another kind of paradigm of computer (Inaudible) communication that's coming about, which I think will again manifest itself in the next twenty-five years, and it's that which we call mobile or nomadic computing. I shouldn't have to sit at my desk to (Inaudible) access to a network (Inaudible) computing. I should be able to go anywhere, in fact, here today I have a laptop with me and I can connect into this facility, get to my files at UCLA, add my E-mail, add new files. Right now that's very cumbersome.

The idea of having information devices min the environment, thousands of small

computers serving me, in my shoes, in my pencil, in the walls, on my furniture, in my personal assistants, in my little computers, things that can display to me and listen to my voice, just the network, the world will be alive with sensors, effectors(?), computers and communications. That's, I think, where we're heading. And the trick is to get the applications right.

We have no idea what the killer(?) application is for the superhighway. No idea. We got to get it right now so that we allow those things to be enabled. The thing not to do now is to create an architecture which basically prohibits applications from arising, which precludes them from arising because of a mistake in the architecture now. So one(?) must have a long range vision at this point. We had that kind of vision back in the early ARPANET days.

We created a structure which was

flexible. Early on, we had the network control program, the NCP, which evolved into the TCP, and the Internet protocol, which (Inaudible) things to come in underneath it. Just lift them up, shove the new thing in, and it works. That's the philosophy we need now to bring about the kind of world I tried to describe.

QUESTION

LEN KLEINROCK

Twenty-five years from now, I would like this world to be alive with intelligent devices. Small computers imbedded in my environment, that I (Inaudible) in a variety of ways, that I can ...

(END TRANSCRIPTION)

(END OF TAPE #6)

Okay. You might wonder who or where, what was really happening in terms of the impact, the long range impact. The answer's absolutely not. For example, BBN structured the machine so it could not support more than sixty-four computers ever. And later (Inaudible) had to expand that. Well, you know, now we have three million computers attached to the network, growing at over eighty percent a year. Nobody had any concept of the impact of this development at all.

You know, there's another story about this ... this IMP that arrived. This B ... DDP-516 was a machine that Honeywell featured in 1968, and that's why many people bid it. Well I remember seeing that machine and I believe it was, either the string(?) or the four joint computer conference in one of the major cities, probably Las Vegas. And they had this machine, it was a hardened(?) military version, hoisted up,

hanging by it's hooks, swinging and running. And there was some son of a gun with a sledge hammer whacking the damn thing, proving that it could continue to operate while they beat the hell out of it.

I believe honestly that that's the physical machine that, (Laughs), BBN delivered to us, and sure enough, it continued to work. And as you know, it's still at UCLA today. Decommissioned, but proudly standing there as an epic(?) of the past.

QUESTION

LEN KLEINROCK

Can I say something mean about them then?

(OFF MIKE)

LEN KLEINROCK

I mean, I'm ready to strangle my secretary too. There's a lot, there were, there were lots of screw up's there.

(OFF MIKE)

QUESTION

LEN KLEINROCK

The technology behind packet switching rests on some very sophisticated mathematical structures. And I hate to say it, but it's ... it's related to something called queueing(?) theory. Now, the word queueing can be spelled one of two ways, and I prefer to spell it Q-U-E-U-E-I-N-G, because you get five vowels in a row, it's the only word in English I know with five vowels in a row.

Queueing theory is a, is a theory which talks about how long we spend waiting in line. Now what the devil can that have to do with networking? Well the answer is this. When these packets, these blocks of data move through the network, they reach a switch and maybe somebody is being transmitted out of that switch when they want to go, so they've got to wait in line. So queues build up, delays get included, and if you listen to the next words I'm

gonna use, you'll see why it's important.

We talk about response time, throughput, blocking, capacity, delays, those are exactly the measures you need to evaluate how a network's performing. When I send a message through a network, I want to know how long is it ... it gonna take to get there? How many bits per second can I pump in? What is the capacity of this network in the first place? So these issues have to be dealt with with a rather sophisticated theory, and queueing theory is exactly the right mathematical structure.

For my own dissertation, I used that theory and basically enhanced it, so it would apply to packet switching. And there's some very sophisticated things that have to be done, but that's why this weird notion of queueing theory enters in the first place. You've got to be able to predict the way a network's gonna perform before you build it. You've got to be

able to decide how to design a network. In order to do so, you have to evaluate way, the way different designs will perform.

And so it's a rather sophisticated mathematical structure. You could do it analytically, you could do it by simulation, you can build the damn thing and then measure it and then it's too late. We chose to do the analytic approach, and the simulation approach.

QUESTION

LEN KLEINROCK

Okay.

QUESTION

LEN KLEINROCK

In 1966, Howie Frank, who was then an assistant professor up at Berkeley, and I was a, an assistant professor down to UCLA, called me and asked me to participate in a short course he was running up there. He had heard of me because I had published a book called

"Communication Networks", that was the name of my thesis. By the way, that made him pretty mad, I understand, cause he had hoped he could use that title for a book he later wrote, and he's constantly complaining that I stole the ... the only right name for a book of this sort.

Now Howie at that time was involved in networking things as well. He was involved with topological design procedures, where to put lines in a network. What kind of a topological structure do you want? I was concerned with the related issues. I would take a problem, for example, if you give me a network with a certain topology, how will it perform? And if you change the topology, how will it perform? Howie was more focused on showing what the topology should look like.

Well. So I was scheduled to give this short course, (Laughs), at Berkeley, but I was plan, I had planned a vacation in, uh, Sequoia,

that would have been the end, so I closed the door, put the belt back on.

I would have rolled it on the right side, but it was a, again, not a stiff wall but it would turn the car over. There I am. What would you do in a case like this? I did the only thing I knew left, I yelled help. (Laughs) And believe it or not, help came. Suddenly my rear left tire blew and it caused the car to slowly grind to a stop. I'll tell you later what happened, why that occurred. But meanwhile, I got to get to Howie. Now I need money to repair the car, and he was paying me some good money to give this lecture.

So I hop out, a park ranger comes by, he says, yep, your car stopped. I said, good. Take me out of here. So he took me to a garage, I told the garage guys, fix the car, and I started

of Sequoia. Ranger comes back

the park, hitchhiking, looking like a bum and nobody's gonna pick me up.

Finally I get a ride, racing down to the airport in ... in ... in, uh, Fresno. Mad rush, run through the airport with my beard, my suit carrier, hop on the plane, just made within a minute. Get on the plane, we taxi out to the runway, plane stops. And the captain says, folks, we have a problem, our brakes aren't working. Takes us back to the airport, hop on any plane that's in Fresno, just to get out of Fresno, and this plane was going to Los Angeles. Hop on the plane, to make a long story short, I got there three hours late. Howie was amazed I got there at all.

(TAPE SPEED VARIES FROM THIS POINT ON)

KLEINROCK

gonna talk the next day. I did the afternoon and I came back a few days later to finish up my half. Got back to my tent in Sequoia that night, and my wife casually says to me, did you have a good day? (Laughs)

QUESTION

LEN KLEINROCK

Okay.

QUESTION

LEN KLEINROCK

Okay. What happened between Howie Frank and myself was we met in this 1966 lecture. Shortly after that, as you know, Larry approached me to put together the ARPANET. And as we began to grow the ARPANET, you know, to connect four nodes together is easy, there's no topological design. When you start

nodes, it becomes a

I introduced him to Larry and Larry hired him on as a sub-contractor, to now take over the topological design of the network. Which then BBN would deploy. Howie's approach, as I said, was topological. Mine was analytic and delay oriented, and we built our different approaches, which tended to lead to the same kinds of design. So Howie, Larry and I worked very well together, it was a very good mix. Bob Kahn(?) was also, of course, making his presence known then, through his efforts at BBN, and as you know, later on, Howie Frank, Bob Kahn and myself wrote this classic paper, which described our experiences in putting the network together.

(OFF MIKE)

(END OF TAPE #5)

BBN JOB #41540

"BBN Job #41540"
Len Kleinrock
Tape Number 6
-- CARIBINER GROUP

QUESTION

LEN KLEINROCK

In 1967, I first became an ARPA principal investigator, and I inherited a staff of programmers, graduate students and researchers who were involved in a variety of computer science efforts. In that group was Vince Surf(?), Steve Crocker(?), John Pastell(?) and Mike Wingfield(?). And a number of other programmers and people. Now, we were doing research on time sharing systems, and that was my specialty, I was doing modeling. And I suddenly inherited forty people and I started managing a large group.

You have to understand, programmers are a special breed. Namely they are totally

uncontrollable. Now I remember those three guys, Surf, Crocker and Pastell, they'd walk into my office, and they would maneuver me and try to get out of me travel to the most exotic places, or equipment or software far beyond our budget. We had this constant tension between us, and very often, I caved in and gave them what they wanted, which was a ... exactly the right thing to do. Cause you have to give people like that their head.

Then, in '67, we began to think about the networking, we had to implement it in '69, and that was the team that I brought together, these young graduate students, brilliant, uncontrollable, full of ideas, and sent them off, in fact, to first write the IMP to host software, and after we got that first IMP working, and then after we got the other four node, the other three nodes, the Stanford Research Institute, University of Santa Barbara, and University of

Utah, the next job was to get computers to talk to each other.

Not just IMP to ... to computer, but computer to computer. In those early days, there was almost no use of the network. In fact, one of the first uses was Rand Corporation in Santa Monica, had been taking decks of cards over to the UCLA large computing machine to run some globic(?) climatology studies. Once we had our switch at UCLA, they could send their data to our switch and the switch was connected to the UCLA machine. One of the first uses was not a network activity, it was what we call incest.

Incest means it comes into one switch and goes out the same, back to another host on the same switch. In fact, one of our early measurements show that twenty-two percent of the ARPANET traffic was incest. I coined that term, if you will.

So we didn't have any good host to host software, and as more ~~Christmas~~ campuses were attached, the main uses of the ARPANET in terms of applications was when someone from one location would take a job at another location, both of which were connected on the ARPANET, and the first guy wanted to use his old machine. So he'd use it through the ARPANET. The ability for two computers to talk to each other was outrageously difficult at that time. And immediately, of course, we set out to build a host-host protocol.

And again, who else to do it, no better than the graduate students, from this community across these few universities. We gave them the job to do over one summer, and it took two years. But finally it got done, the host-host protocol was established, it began to get implemented in various hosts, and then activity began. And sometime around then is when

somebody had the idea, listen, why not do this thing called electronic mail? We will ad-hoc add on. Of course, now that dominates network traffic. It was not at all anticipated when it wa first used.

We're now into the early Seventies. We have cross country communications, twenty-odd nodes attached, it's time to announce the network to the world. And so ARPA organized to attach themselves to the International Conference on Computer Communications, in 1972, October, at the Hilton in Washington. Now Washington was not in the ARPANET at the time, and they wanted to run a demonstration of the ARPANET at that conference.

So Larry managed to order lines, get a switch in Washington, and provide the capability. And he announced to the community that each of his contractor, that he was supporting for all these millions of dollars, would have to provide

some application. And write up the application in a book, and the idea, passers-by at that conference would be dragged into this exhibit, given a book, and said, use this thing called a network. And then we'd get somebody to help them use it.

So of course, everybody prepared these great experiments. Lots and lots of applications on ... on ... on various time sharing computers. A lot of chess playing programs, a lot of artificial intelligence, and we had concocted the following experiment. The idea was, you sit down at a machine in Washington, at a teletype, and you log on to a machine at UCLA across the country. Have that machine pull up a file from the machine in ... at BBN in Boston, ship that file back to the machine at UCLA, and then print it out in Washington. Very nice demonstration of the networking.

The night before the exhibit began,

(Laughs), I remember this very clearly, I believe it was John Pastell was at the keyboard, and he logged onto UCLA, pulled up the file from Cambridge, shipped it to his, print to right next to him, and nothing happened. Nothing happened. And he looked around the room, and there were many other demonstrations there, one of them was, MIT's Turtle. Turtle is a little robot that runs around the room. Turtle was jumping up and down.

The system had accidentally sent the data to Turtle. (Laughs) And Turtle was interpreting it as motion and jumping around. They fixed that, it was a great demonstration once they finally got it going.

(OFF MIKE)

(CUT)

QUESTION

LEN KLEINROCK

At that exhibit in 1972 when we were first

giving the first public demonstration of the ARPANET, the community at large was aware of the ARPANET, but not at all familiar with what it could do. So a number of phenomena. First of all, in spring 1972, there was a Spring Joint Computer Conference. A group of us had published papers describing various aspects of the ARPANET. I believe it was a set of four papers which became a set of classics. I had a paper, Larry had a paper there, there were a number of papers there.

So the science community, the computer community, was aware of the ARPANET. Here was their first time to try to kick the tires. But something else happened. The ARPANET community was also challenged to make this a successful demonstration. so people were busy writing these wonderful interesting applications. And when people came to that exhibit, they were totally overwhelmed by

the capacity of this network to perform functions.

Here's another experiment, another demonstration. It was a simulation of air traffic control. The idea was to simulate an aircraft taking off at Boston and landing it in New York. So they had a simulation running on a machine physically in Boston, simulating air traffic control, and this simulated aircraft now leaves the Boston region and enters the Hartford, Connecticut region. A second machine picks up the air traffic control for ... for Hartford, and when it passes to New York, a third machine picked that up. Actually, three different computers were sharing this simulation, moving data from one to the other.

A very effective demonstration again. It was a huge success. Uh, people were overwhelmed, they used it very easily, lots and lots of artificial intelligence applications and

networking applications.

QUESTION

LEN KLEINROCK

The sense in that room was not one of fear or concern, it was excitement. I mean, here we could show it off, we knew it would work. Even if it fumbled, I mean, these things were fixable. It was a wonderfully exciting experience, and in fact, it had a ... a magnificent step up in the use of the ARPA network. We were tracking how much traffic was in the ARPA network month by month. In fact, BBN was doing the monitoring of the usual traffic, and at UCLA, we were the ARPA network measurement center.

We were responsible for experimenting with the network, it was our job to break the network? We could break it on call. (Clicking Sound) In fact, there were lots and lots of deadlocks that came out once we started testing that (Inaudible). But going back to this

measurement of the traffic. You could see that the traffic in October, 1972 jumped and it stayed high for a period of time, and the reason is, we created many, many guest accounts for that demonstration. And they remained alive for a period of a few months. So it totally increased the traffic artificially at that point.

QUESTION

LEN KLEINROCK

Don't have any good jokes, but I could tell you about some good deadlocks.

QUESTION

LEN KLEINROCK

Okay. Our job as network measurement center, as I said, was to experiment with the network. Namely, to stress it, create conditions which would test the envelope of network behavior. And we were doing that by enabling measurement activity throughout the network, we could generate art ... artificial traffic, et

"BBN Job #41540"
Len Kleinrock
Tape Number 5
-- CARIBINER GROUP

(OFF MIKE)

QUESTION

LEN KLEINROCK

First time I heard about the ARPANET was when Larry Roberts(?) approached me and said, Len, I've got this great idea and this great need. We've got to connect computers together through a network for a variety of important reasons. His main reason was that he wouldn't have to buy a new computer for every researcher that he supported in ARPA(?). So the original idea of the ARPANET was to do this sharing of resources. And it wasn't a military application that he told us.

I don't know what he told the Pentagon, but his ... his ... his idea was, he wouldn't have

to buy everybody a computer, and those that he did buy computers were doing special things with them. Example, at Utah they were doing special graphics, UCLA we were doing simulation packages, uh, to SRI(?) they were doing information and storage and retrieval. And these became very unique resources, there was a super computer being developed at Illinois.

M — The specialized things everyone was doing...

So he said, look. These unique resources need to be shared over the network, and I'm tired of buying every new re, investigator a new computer, so let's connect them together. That was the first knowledge that there was gonna be a thing called the ARPANET created.

(OFF MIKE)

(CUT)

(OFF MIKE)

QUESTION

LEN KLEINROCK

batteries, no power, no nothing, it was free, and in some sense, an engineer was born that day.

Spent the rest of my youth cannibalizing discarded radios, putting together radios, designing them, building them. Went to Bronx Science, took the usual courses and studied radio engineering, and then I went off to City College. Couldn't afford to go, by the way, to a day session. Went to evening session, got a bachelor's degree in electrical engineering, was elected student body president of the evening session and got a scholarship to, full scholarship to MIT.

And that's when this networking side of my life begins, with the strong electrical engineering background.

QUESTION

LEN KLEINROCK

Not ...

BBN JOB #41540

BBN JOB #41

BBN

LEN KLEINROCK

At MTI, some interesting things were going on. The first thing you do as a Ph.D. student ...

QUESTION

LEN KLEINROCK

Oh.

QUESTION

LEN KLEINROCK

In 1959, I began a Ph.D. program in the electrical engineering department. Now all of my class mates were busy refining a highly specialized area in communication theory, known as coding theory. And I decided that's not for me, it were just too little, too late. I wanted to break new ground. And so I began to work on this thing called a message switching communications network, and I began to study what later would become the notion of packet switching.

I studied the principles, the ar
design, mathematical treatment

BBN JOB #41540

BBN JOB #41r

B

cetera. And around that time, I met Larry Roberts, he was a classmate of mine. He was doing something totally different, he was working on three-dimensional, uh, picture processing. And I also met Ivan Sutherland(?) at that time, a name that may or may not be ... be known to this group. Um, a key guy.

Anyway, Larry and I were strong classmates, and I was doing my work on networking and he was doing his work on picture processing. And, uh, I graduated, my work was finished in December '62, it became a book, by McGraw-Hill, Lincoln Laboratories, MIT, in 1964, and laid down the foundations for what later became the packet switching technology. Unfortunately, there was absolutely no one interested in this technology in 1962, '6

In '63 I joined the faculty

continued to pursue the res

luatic

BBN JOB #41

BBN JOB #41540

then in 1967, Larry approached me and said, Len, I've got this need. Every time, he was now working for ARPA, every time a new investigator wants to be funded, he asks me to buy them a computer. And I'm getting sick of buying, (Laughs), all these computers cause they're all the same machine. And every time they get one of their machines, they make something different out of it.

One guy puts graphics on it, you put a lot of simulation on it. They were doing artificial intelligence at MIT and Stanford. And now these unique resources are out there and I can't duplicate all that capability at every location. And I know you worked on networking. So let's put together a network which allows these various computers to be shared through *' network, so you could use my sim package and I can use your

BBN JOB #41540

BBN JOB #41F

So in 1967, Larry gathered around him a small handful of people, including myself, to lay out the specification for what was later to become the ARPANET. I remember very clearly there was, um, Herb Baskin(?) was there, he had done some work in time sharing. And he banged his fist on the table and said, if this network can't give me a one second response time, it's no good for time sharing. And so we put down a spec, (Claps Hands), one second response time.

I was there, and I said, look, this is an experimental network. If you can't measure what's going on, you've got nothing. So we installed sophisticated requirements for measurement hooks. And so the, slowly, the spec began to evolve. Spec was written in m' 1968, a request for quotation(?) was se industry, a number of industry or responded, BBN being one

BBN JOB #41540

What they were gonna do is, take a mini-computer, change some of the hardware and the software, make it into a machine that could do packet switching in a network. BBN won the contract in nineteen-seven ... sixty-nine, January of '69. And they were supposed to deliver to UCLA the first IMP, Interface Message Processor, the day after Labor Day in 1969. They chose UCLA as the first site because of my strong development of the early technology and my work all those years. That was the game plan.

And our job at UCLA was to be able to receive this switch and connect it to our host computer. Now that made us the first node(?) on the network. You might laugh, how can ' be the first node on a network? And tᵇ is, you got to make those two m each other. When(?) we ʰ

BBN JOB #41540

BBN JOB #415

to our delight, cause we needed that time to get the spec ready(?). In fact, BBN wouldn't tell us how to talk to their machine, and yet, when they delivered the machine they expected us to be able to communicate.

So we had to squeeze out of them, plead out of them, the spec, invent part of the spec ourselves, write the software, build the hardware, et cetera. Machine's gonna be two weeks late, what could be better? Except they then ... then informed us they put the damn thing on an airplane, and it was gonna arrive on time. We worked day and night, we were ready when it arrived, it came in on a Tuesday, the day after Labor Day. You got a picture of the scene, okay? My guys were there. I had put toget' this team ...

QUESTION

''ROCK

BBN JOB #41540

BBN JOB #41540

(OFF MIKE)

(CUT)

(OFF MIKE)

QUESTION

LEN KLEINROCK

The early ideas of packet switching came from a number of sources independently. I in my Ph.D. dissertation, had laid down the mathematical principles, the design, the protocol evaluation, et cetera. The same time, Paul Barron(?) at Rand Corporation, had written some, was writing some reports in the early Sixties about a thing which amounted to packet switching as well. We both had the same idea of creating a mesh network, with dynamic routing through a network, with addressed packets moving through, highly reliable kind of syste'

I did the analysis, at t'

was inventing similar ide

National Physical '

BBN JOB #41540

was New, was ... wasn't Newcastle, it was, remind me of the city somebody. Um ...

(OFF MIKE)

QUESTION

LEN KLEINROCK

The origins of packet switching came from more than one source. In my early work from '59 to '62, I had laid out the mathematical foundation for the way packet switching networks would behave and should be designed. There were mathematical models, simulation models, protocol models. In the early Sixties, Paul Barron produced some reports which described similar ideas from a, from a, uh, behavioral point of view.

He talked about the idea of random routing, as I had. He talked about packets, he talked about dynamic routing, he talked about mesh networks as ... as had I. Meanwhile, Donald Davies(?) at National Physical Laboratory

in England, was thinking about the same things, and he reached a stage of actually building a packet switch, even before the ARPANET was developed. But he never got beyond the one switch. Basically, they ran out of funding in England.

The same time, Larry Roberts, who had been a classmate of mine, in the mid-Sixties, joined the ARPA office, and he earlier had tried to make a connection from a machine at Lincoln Laboratory, to a machine at System Development Corporation, that's from Cambridge to Santa Monica. And he found that so painful, that he marked it in his head, some day he's gonna try to fix that. So here we have in the early Sixties, these pioneers developing technology of packet switching on paper.

Larry's at the ARPA office, and what is he doing, he's supporting research in time sharing. So along comes a new principal

investigator to him, Larry says listen, I'll support you. The investigator says terrific, buy me a computer, if you want me to do research on computers. Now, this was getting a little bit boring for Larry and expensive. He couldn't afford to give everybody a computer. So that, with his knowledge that there was the ability to create a network where these machines could be accessed remotely, led him to the concept of an ARPANET.

(OFF MIKE)

QUESTION

LEN KLEINROCK

Let's talk about where the concept of packet switching came from. In fact, there were a number of independent studies that led to this idea of packet switching. In '59, I began studying the principles, the mathematical analysis, the simulation and design technology for what became packet switching, it was in the

form of message switching at the time. And in '62, I presented my Ph.D. dissertation, which became a book published by, uh, MIT-Lincoln Laboratories.

The same time in the mid, the early Sixties, Paul Barron at Rand Corporation was conceiving similar ideas. These ideas between us were the idea of a mesh network with packets, blocks of data moving through a network independently, with dynamic routing, going around blocked and congested regions, et cetera. Highly robust, highly dynamic, very effective technology.

Meanwhile, in England, Donald Davies was conceiving of the notion of packet switching as well. And he went so far as to build a packet switch, a single note(?) packet switch before the ARPANET actually came on line. Unfortunately, he ran out of funding, and they couldn't pursue that project. But here we have these ideas

coming out of a number of different independent sources.

Meanwhile, I had met Larry at ... at MIT, he and I were very close classmates, in fact, we shared an office at Lincoln Laboratory together. Larry in the mid-Sixties, joined the ARPA office when Robert Taylor, who had been funding computer research in the Sixties, brought Larry in to put together the notion of a network. Bob Taylor basically sold the idea to the government, and Larry was there to make it happen.

Now, why did they want a network like this? Well, in the mid-Sixties the ARPA computer office was supporting the hottest topic around then, which was time sharing. And they went to a number of principal investigators and said listen, you want to do research on time sharing? The guy said yes. If you're gonna make me do research, buy me a computer. And

they kept buying computer after computer, and at each site, the computer changed form.

For example, at UCLA, we, and ... at UCLA, we were able to add very sophisticated simulation capabilities on our machine. At Utah they were adding graphics capability. At Carnegie-Mellon, MIT, Stanford, they were adding artificial intelligence. At Illinois, they were building a super computer. So all these unique resources were appearing throughout the network, and every time a new research ... researcher came on, there's no way that the ARPA office could bring all this capability to each site.

So Larry, having experimented with networking earlier, he had tried to connect a machine from Lincoln Laboratory in Cambridge, Mass., to a machine in Santa Monica at the System Development Corporation. And that was such an unpleasant, difficult procedure for him,

he noted in his head that some day, he'd like to see that done in a more sophisticated way.

So here we have, Larry who knows my work in networking, he recognized there's a need for networking sometime. He suddenly finds a need at the ARPA office, they're spending too much money on providing everybody with the same kind of computer. So let's create a network and make all these resources accessible to people through the network. Fine. (Claps Hands) Larry approached me in 1967 and he said, Len, join me with a few other people and let's write down the specifications to what a packet switching network should look like.

So we went there, and we laid out the spec. I remember very well Herb Baskin was there, and he had done some work on time sharing. And he pounded his fist on the table and said, look, if that network can't deliver a response time of one second, I can't use it for

time sharing, it's too slow. So we wrote down a spec, one second response time.

And I was there and I said, look, if you can't put measurement hooks(?) into this experimental network you'll never know what's going on, so we specified a very sophisticated measurement package, which indeed found it's way into the software of the network. So the spec evolved. In 1968 it was formalized and sent out for request for quotation, a number of organizations responded. And many of them responded with the same machine.

See, the idea of this proposal was to take an existing mini-computer, add some hardware, add some software, and make it behave like a packet switch. Many of these organizations selected the Honeywell DDP-516(?) mini-computer, and BBN was one of those. BBN won the contract. In January of '69 it was awarded to them. And their task was to deliver

the first IMP at UCLA the day before Labor Day in 1969. Turned out to be a Saturday before the Monday of Labor Day.

Meanwhile at UCLA it was our job to build the hardware and the software, to allow our host computer and their IMP to talk to each other. Now that sounds like a straightforward process. That makes us the first node on the ARPANET, which sounds like an oxymoron, how can you be the first node on the network? And the answer is, to get these two machines to talk to each other, was a really sophisticated task. Especially so because BBN would not release the IMP to host specification to us.

And we had to squeeze blood to get it out of them, and in some cases we couldn't and we invented our own specification. So finally we got this back, we were busy trying to implement. Recognizing that the switch is going to appear on a Saturday just before Labor Day. In mid-

August, we learned that BBN had slipped the schedule and they were behind two weeks and we were totally thrilled with that. We needed the time to continue our implementation.

We learned just before Labor Day that the sons of guns, they put the damn thing on an airplane and were shipping it out to us. That kept us up, us up a few nights, finishing up the specification, getting implementation. The machine arrived, and we were ready. Now the scene when that machine arrived is something to behold. Everybody who could be there was there.

First of all, me and my team were there. We had the people from the UCLA administration there. I mean, after all, you're gonna use our university? BBN was there. ARPA was there. Of course Honeywell was there. The people who built my host machine, the system, um, Scientific Data Systems, was there. We had

ATT Long Lines there, we were gonna use their lines. We had GTE, we were gonna use their local access. Everybody that could be there was there, and they're all ready to point the accusing finger at the other guy, if it didn't work.

And fortunately, bits were flowing that first day. the interface was a success. And by the following Tuesday, we had messages moving back and forth. Fortunately, it was a big success and that was the birth of the ARPANET. It happened right then, around Labor Day, 1969. It was a great success.

QUESTION

LEN KLEINROCK

Well I feel very stiff, frankly, but ...

(OFF MIKE)

(CUT)

(OFF MIKE)

QUESTION

LEN KLEINROCK